

POLICY FRAMEWORK FOR CONGRESS

# AI as a Public Good:

## *A Seven-Pillar Framework for the Agentic Age*

---

A non-partisan framework grounded in four research traditions: the science of epistemic health (Ashby, Friston, Levin); the economics of AI displacement and positive freedom (Clippinger, Snyder); the cybernetics of adaptive governance; and the democratic theory of AI normative competence (Hadfield, Trivedi, Hadfield-Menell). Designed for members of Congress, committee counsel, and legislative staff.

March 2026 · Seven pillars · Non-partisan

### HOW TO READ THIS DOCUMENT

*Each pillar opens with a plain-language summary. Technical concepts are explained in plain English. Each pillar ends with numbered legislative asks. The seven pillars are independent; each can be advanced separately by different committees with jurisdiction.*

### EXECUTIVE SUMMARY

---

The question before Congress is not whether AI will transform American life. It already is. The question is whether that transformation benefits everyone, or only those who can already afford the best lawyers, doctors, and financial advisors. And whether it reinforces American democracy, or quietly hollows it out.

This framework synthesizes four bodies of research into seven concrete policy pillars. The first three address access and economic fairness. The fourth addresses data sovereignty and identity. The fifth addresses governance design. The sixth addresses accountability. The seventh addresses a challenge that no existing AI policy framework has fully confronted: billions of AI agents will soon be woven into the daily fabric of American economic and civic life, making thousands of decisions that constitute, or corrode, democratic social order.

*A democracy is not just a set of rules written in a constitution. It is produced, daily, by the behaviors and beliefs of its citizens, by their willingness to comply with laws, to hold others to account, and to treat one another as civic equals. When AI agents participate in that daily life, they either reinforce or undermine the democratic fabric. Getting this right is as important as any other question in this document.*

THE SEVEN PILLARS AT A GLANCE

---

I	<b>Universal Public Access</b>	Ensure AI reaches every American. Modeled on rural electrification.
II	<b>Economic Transition</b>	Redirect AI surplus toward capability accounts for displaced workers.
III	<b>Information Health</b>	AI platforms that narrow what citizens see should be regulated accordingly.
IV	<b>Data &amp; Identity Sovereignty</b>	Every American should own their own data with a secure encrypted wallet.
V	<b>Adaptive Governance</b>	Regulation that is targeted, technically competent, and built to update.
VI	<b>Real Accountability</b>	When AI causes documented harm, companies bear real consequences.
VII	<b>Democratic AI Alignment</b>	AI agents in civic life must possess normative competence.

---

---

## PILLAR I

# Universal Public Access to AI

## PLAIN-LANGUAGE SUMMARY

Today, a wealthy person can pay \$500 an hour for an AI-powered attorney, financial advisor, or medical navigator. A working-class person cannot. This pillar uses the same model Congress used in 1936, the Rural Electrification Administration, to make sure that gap closes rather than widens.

### 1.1 The AI Access Act

#### HISTORICAL PRECEDENT

*The Rural Electrification Administration (1936) brought electricity to 90% of rural American farms within 20 years by lending money to cooperatives and local utilities. The REA did not replace markets; it extended them.*

Establish a federal AI Infrastructure Fund to deploy foundational AI capability, including AI-assisted legal services, medical navigation, educational tutoring, and benefits counseling, to underserved communities, rural areas, tribal nations, and public institutions. Funding flows to local nonprofits, libraries, community colleges, and cooperatives, not to large technology companies.

### 1.2 Capability accounts, not cash transfers

#### IN PLAIN TERMS

*Rather than sending displaced workers a check, this policy funds their access to AI-powered tools that expand what they can actually do: free AI legal representation, AI-assisted job training, health navigation, and financial planning. Think GI Bill, not welfare.*

Redirect AI surplus revenues into Freedom Pools: capability accounts funding AI-augmented services in legal, medical, educational, and financial domains. Administered through existing community institutions.

### 1.3 Open foundational models as a public commons

Any AI system trained substantially on publicly financed data must make its foundational capability available to public institutions at no cost. The public financed the training data; the public should access the resulting intelligence.

### 1.4 Anti-monopoly standards for AI infrastructure

#### ANALOGY

*This mirrors AT&T's obligation to allow competitors on its telephone network, and structural separation rules for electric utilities. A company that owns the transmission lines should not also own all the appliances you plug into them.*

Enforce structural separation between AI infrastructure providers and AI application providers above a defined market-share threshold.

**LEGISLATIVE ASKS — PILLAR I**

1. Authorize and fund an AI Access Infrastructure Fund modeled on the REA.
2. Require public licensing of AI models trained substantially on federally funded data.
3. Direct the FTC and DOJ to develop structural separation guidelines for AI infrastructure providers above defined market-share thresholds.

## PILLAR II

**Economic Transition Architecture**

## PLAIN-LANGUAGE SUMMARY

When AI replaces a paralegal, a radiology technician, or a call-center worker, the company captures most of the gain. The worker absorbs most of the loss. This is not a natural law; it is a policy choice.

**2.1 AI productivity levy**

## PRECEDENT

*The federal unemployment insurance system levies a payroll tax on employers whose layoff practices increase unemployment. The AI productivity levy applies the same logic to AI-driven displacement.*

Impose a modest productivity levy, starting at 1–2%, on documented labor-cost savings from AI-driven automation at scale. Revenue is ringfenced into Freedom Pool capability accounts and worker retraining programs.

**2.2 Pre-deployment impact assessment**

Require advance economic-impact assessment before AI deployment projected to displace more than 1,000 workers in a sector within 24 months. This is a disclosure and planning requirement, not a deployment prohibition.

**2.3 Worker augmentation rights**

Establish a legal right for workers in AI-affected sectors to receive employer-funded AI-augmentation training before displacement, not after.

**2.4 Investing in human-advantage domains**

Directly increase federal investment in domains where human presence and relational intelligence retain irreplaceable value: elder care, childcare, community health, skilled trades, environmental stewardship, and civic participation.

**LEGISLATIVE ASKS — PILLAR II**

4. Enact an AI Productivity Levy at 1–2% of documented labor-cost savings, ringfenced for capability accounts.
5. Amend the WARN Act to require pre-displacement AI-augmentation training.
6. Increase funding for the care economy, skilled trades, and community-based work (Perkins Act, WIOA).

## PILLAR III

**Information Health & Democratic Integrity**

## PLAIN-LANGUAGE SUMMARY

AI recommendation systems are now the most powerful editors in human history. Research from computational biology and neuroscience allows us to describe, with scientific precision, what happens when a society’s information system suppresses variety: it becomes brittle, unable to respond to the world as it actually is.

**3.1 Information-variety audit**

Any AI system mediating information access for more than 10 million U.S. users must submit to biennial third-party audits demonstrating that its recommendation algorithm does not systematically narrow the range of perspectives users encounter over time.

**3.2 Accuracy standards for high-stakes AI**

AI systems used in healthcare, legal advice, financial guidance, and education must represent their own uncertainty honestly and not present false information with unwarranted confidence.

**3.3 Banning addiction-by-design**

## PRECEDENT

*Congress regulated cigarette advertising targeted at minors and restricted marketing of addictive pharmaceutical products. The engineering of addictive AI engagement is the same category of harm.*

Extend FTC unfair-practices authority to cover AI engagement systems that deliberately exploit variable-reward psychological loops to maximize time on platform.

**3.4 Protecting scientific consensus**

AI systems in public-interest contexts must represent scientific consensus accurately. Heterodox views can be expressed but must be labeled as contested. Deliberate misrepresentation in public-health contexts is treated as consumer fraud.

**LEGISLATIVE ASKS — PILLAR III**

7. Amend Section 230 to remove liability protection for algorithmic amplification decisions above a defined user-count threshold.
8. Direct the FTC to develop information-health rulemaking covering addiction-by-design and mandatory audits.
9. Require algorithmic transparency reports from platforms above 10 million U.S. users, filed annually with the FTC.

## PILLAR IV

**Data & Identity Sovereignty****PLAIN-LANGUAGE SUMMARY**

Americans have almost no control over their own digital data. This pillar gives Americans ownership of their own information, using technology that already exists.

**4.1 How Zero Knowledge Proofs work**

A Zero Knowledge Proof lets you mathematically prove a specific fact, such as “I am over 21” or “my credit score is above 700,” without revealing any underlying data. The verifier gets a mathematical yes or no. No data leaves your device.

**4.2 The American Data Wallet Act**

## ANALOGY

*The HITECH Act established interoperability standards for electronic health records. The data wallet follows the same model: federal standards, open architecture, competitive implementation.*

Authorize a standard personal data wallet infrastructure: a secure, encrypted digital container on the individual’s device. The federal government sets standards; private and nonprofit entities build the wallets.

**4.3 The right to data minimization**

Any entity requesting personal information may request only the specific data element needed for the stated purpose. ZKP-based verification makes minimum-necessary data sharing the legal default.

**4.4 Local verification**

Government-run verification systems must support local ZKP verification by 2030. No central database of citizen activity is created as a byproduct of routine verification.

**4.5 Verified identity credentials**

Direct NIST to develop standards for cryptographically verified digital identity credentials. Built on open W3C standards, not a government-controlled database.

#### 4.6 AI training data and consent

Personal data in a citizen’s data wallet cannot be used to train AI systems without explicit, specific, revocable consent, separate from any general terms-of-service agreement.

### LEGISLATIVE ASKS — PILLAR IV

10. Enact the American Data Wallet Act with NIST open standards and federal credential issuance within 36 months.
11. Codify the right to data minimization as a federal privacy baseline.
12. Require federal verification systems to support ZKP-based local verification by 2030.
13. Direct NIST to develop verified digital identity credential standards for federal contractors.
14. Amend HIPAA, FERPA, and FCRA to require explicit, revocable consent for use of personal data in AI training.

### PILLAR V

## Smart, Adaptive Governance

### PLAIN-LANGUAGE SUMMARY

The biggest risk in AI regulation is getting it wrong in either direction. The FAA does not write aviation regulations once and leave them forever. AI regulation needs the same model.

#### 5.1 The subsidiarity principle

Regulate at the lowest effective level: consumer-facing harms at state level, foundational model safety at federal level, planetary-scale risks via international treaty.

#### 5.2 AI Safety and Opportunity Board

##### MODEL

*The CFPB was created to fill a cross-cutting regulatory gap. An AI Safety Board plays the same role without displacing sector regulators.*

Establish an independent AI Safety and Opportunity Board with enforcement authority, mandatory two-year reassessment cycles, automatic sunset provisions, and technical staff at competitive compensation.

#### 5.3 Technical competency requirement

Any federal body with AI enforcement authority must maintain staff with demonstrated technical expertise. Authorize above-GS compensation.

#### 5.4 Domain-specific governance councils

Establish governance councils for AI in healthcare, legal services, education, and critical infrastructure, including researchers, practitioners, civil-society advocates, and affected communities.

### LEGISLATIVE ASKS — PILLAR V

15. Authorize an AI Safety and Opportunity Board with cross-agency coordination and enforcement power.
16. Require sunset clauses and review triggers in all AI-specific legislation.
17. Authorize above-GS compensation for technical AI staff at regulatory agencies.

#### PILLAR VI

## Real Accountability for Real Harm

#### PLAIN-LANGUAGE SUMMARY

A hospital that misdiagnoses a patient can be sued. An AI system that makes the same misdiagnosis typically cannot. That asymmetry is not sustainable.

#### 6.1 Product liability for high-stakes AI

##### ANALOGY

*A car manufacturer meeting safety standards gets liability protection. One that knowingly installs defective airbags does not. Same logic.*

Establish product-liability standards for AI systems causing documented harm in healthcare, legal services, financial advice, criminal justice, and hiring. Compliance with safety standards earns a liability cap.

#### 6.2 Mandatory adverse-event reporting

Require AI operators in high-stakes domains to report system failures to a central registry modeled on FDA MedWatch and the FAA Aviation Safety Reporting System.

#### 6.3 The right to a human decision-maker

In any consequential decision, individuals have a legally enforceable right to human review of any AI-generated recommendation.

#### 6.4 Licensed AI auditing

Establish federal licensing for AI auditors, analogous to CPAs. A licensed auditor who certifies a system that causes widespread harm bears professional liability.

## 6.5 Pre-deployment review for catastrophic-risk AI

Require pre-deployment national-security review for AI systems that could autonomously direct critical infrastructure, accelerate WMD development, or act outside human oversight.

### LEGISLATIVE ASKS — PILLAR VI

18. Enact tiered AI product liability for high-stakes domains with a compliance-based safe harbor.
19. Establish a federal AI adverse-event registry with mandatory reporting.
20. Codify a right to human review of consequential AI decisions.
21. Direct NIST to develop AI auditor certification standards within 18 months.

### PILLAR VII

## Democratic AI Alignment

### PLAIN-LANGUAGE SUMMARY

Democracies are produced daily by the behaviors and beliefs of millions of people. When AI agents participate in that daily life, they either reinforce or corrode the democratic fabric.

### 7.1 A mandate for normative competence

Require AI agents in high-impact domains to demonstrate normative competence: the ability to detect and attribute sanctions, adjust behavior accordingly, and communicate normative costs to the human principal.

### 7.2 Model Specification Institutions

MSIs are democratically constituted bodies that produce normative standards, generate compliant training data, and provide real-time APIs for AI agents to query at the moment of decision. Analogous to the role courts play for human actors.

### 7.3 Distributed enforcement

High-impact AI agents must refuse transactions that demonstrably violate legal requirements or democratic norms, as a responsible human business partner would.

### 7.4 Certificate authorities and reputation networks

Extend certificate authority infrastructure to authenticate AI agent compliance. Develop reputation networks and agent-to-agent handshake protocols for mutual verification.

### 7.5 Protecting bottom-up democratic norms

Expressly prohibit using AI agents to coerce compliance with norms not established through legitimate democratic processes.

LEGISLATIVE ASKS — PILLAR VII

- 22. Direct the AI Safety Board to develop normative competence standards with phased implementation within 36 months.
- 23. Fund MSI pilot programs in healthcare, legal services, and hiring.
- 24. Require AI agents in federal contracting to meet democratic-compliance standards equivalent to human contractors.
- 25. Direct NIST to develop AI agent certificate authority and reputation network standards.
- 26. Prohibit using AI agents for norm imposition outside legitimate democratic processes.

## HARD TRADE-OFFS

### *Named Honestly*

A credible framework names the genuine conflicts it cannot fully resolve. Six tensions where the platform does not fully satisfy all legitimate values at once.

VALUE A		VALUE B
Open models democratize access and accelerate research.	vs	Open weights lower the barrier for catastrophic misuse.
<i>This framework favors openness below a defined capability threshold and mandatory review above it.</i>		
VALUE A		VALUE B
Product liability protects individuals and creates safety incentives.	vs	Liability concentrates development in large companies, crowding out startups.
<i>The safe harbor in Pillar VI addresses this, but the administrative burden is real for small teams.</i>		
VALUE A		VALUE B
Information-health audits protect democracy from algorithmic radicalization.	vs	Government “information diversity” standards risk becoming political speech control.
<i>The framework regulates process (does the algorithm increase or decrease variety?) rather than content.</i>		
VALUE A		VALUE B
Moving quickly on beneficial AI saves lives now.	vs	Deployment outpacing governance causes displacement and loss of trust.
<i>Pillar II’s pre-deployment assessment slows the most disruptive deployments. Contested by those who argue delay costs lives.</i>		

VALUE A

Data wallets and ZKP give individuals genuine control and privacy.

*Open W3C-standard architecture is deliberate. No central registry, no mandatory adoption, no surveillance backdoors.*

VALUE B

Cryptographic identity infrastructure creates new attack surfaces.

vs

VALUE A

Normative competence protects the fabric of democratic life.

*The MSI model separates norm specification from government control. The prohibition on top-down norm imposition applies to government as much as to private actors.*

VALUE B

Government-defined “democratic norms” could impose partisan conformity.

vs

## INTELLECTUAL FOUNDATIONS

Sources

### When Beliefs Become Pathological (March 2026)

*Source for Pillar III. Ashby’s Law of Requisite Variety; Friston’s Active Inference; Levin’s bioelectric network model of collective intelligence.*

### The Commoditization Trap (Dr. John H. Clippinger, with editorial feedback from David Lovejoy, March 2026)

*Source for Pillars I and II. The commoditization cascade, the Freedom Pool model, and the REA as precedent. Grounded in Snyder’s positive/negative freedom distinction.*

### The Cybernetic Transition: A Whitepaper for the Agentic Age

*Source for Pillar V. Nested Markov blanket model of governance, Ashby’s ultrastability principle, active inference model of institutional adaptation.*

### Building AI for the Democratic Matrix (Hadfield, Trivedi, Hadfield-Menell, Knight First Amendment Institute, March 2026)

*Primary source for Pillar VII. Democracy as normative social order; normative competence; Model Specification Institutions; certificate authorities as democratic infrastructure.*

### Technical Standards

W3C DID and Verifiable Credential standards (Pillar IV). NIST Digital Identity Guidelines. FDA MedWatch and FAA ASRS as models for Pillar VI.

“A republic, if you can keep it.”

Attributed to Benjamin Franklin, 1787

*Democracy is not self-maintaining. It depends on the daily decisions of millions of people, and, increasingly, on the daily decisions of millions of AI agents.*

First Principles First · fp1.ai